Distributed Compressed Sensing for Sensor Networks with Packet Erasures

Christopher Lindberg, Alexandre Graell i Amat, Henk Wymeersch
Department of Signals and Systems, Chalmers University of Technology, Gothenburg, Sweden,
Email: {chrlin,alexandre.graell,henkw}@chalmers.se

arXiv:1404.3078v2 [cs.IT] 31 Mar 2015

***Abstract*—We study two approaches to distributed compressed sensing for in-network data compression and signal reconstruction at a sink. Communication to the sink is considered to be bandwidth-constrained due to the large number of devices. By using distributed compressed sensing for compression of the data in the network, the communication cost (bandwidth usage) to the sink can be decreased at the expense of delay induced by the local communication. We investigate the relation between cost and delay given a certain reconstruction performance requirement when using basis pursuit denoising for reconstruction. Moreover, we analyze and compare the performance degradation due to erased packets sent to the sink.**

I. Introduction

Wireless sensor networks (WSNs) provide a tool to accurately monitor physical phenomena over large areas [1]. The WSN is usually considered to be energy-constrained and to comprise up to several thousands of nodes. However, smart phones and other sensing devices carrying powerful batteries have become ubiquitous. This provides a possible platform for WSNs where energy is not a scarce resource. Instead, the sheer number of sensors puts a strain on the bandwidth available for communication between the sensors and the sink. Consequently, the measurement data acquired by the sensors needs to be compressed. Compression should be able to operate under unreliable communication conditions and be scalable in the number of sensors. Existing techniques, such as Slepian-Wolf coding and transform coding (see [2] and references therein) require precise statistical models about the phenomena of interest. Compressed sensing (CS) [3]–[5], on the other hand, alleviates the need for precise statistical models and is scalable in the number of sensors [6].

Prior work on CS in WSNs includes [7]–[11]. In [7], [8], CS is used for in-network compression, but communication to the sink is done by analog phase-coherent transmissions. This is not practical for WSNs operating in a cellular network since all sensors need to be perfectly synchronized. In [9] and [10], CS is considered for networks with multi-hop routing towards the sink. In addition, [10] considers the delay caused by a medium access control (MAC) protocol. The drawback of multi-hop routing is the necessity to form a spanning tree, which is impractical and prone to communication failures, especially when the sensors are mobile. In [11], no sink is present, but the sensors use CS and consensus to perform distributed signal reconstruction. However, the focus is on reconstruction performance and the MAC delay is not studied.

In this paper, we consider distributed CS for a WSN with equispaced sensors on a straight line. The sensors sense a physical phenomenon in their local environment, perform in-network compression, and transmit the (compressed) data to a common sink. We analyze the tradeoff between communication cost towards the sink and MAC delay from the inter-sensor communication. We consider two approaches that rely on local processing between sensors, where only a subset of the nodes communicate to the sink. The first approach performs local processing by clustering of the sensors, while the other uses average consensus. Additionally, we compare the robustness to packet erasures when transmission to the sink is performed over a noisy (erasure) channel. Our contributions are:

- Closed-form expressions for the upper bound on the reconstruction error for basis pursuit denoising (BPDN), that guarantees stable reconstruction for both approaches in the presence of packet erasures.
- Closed-form expressions for the communication cost and the MAC delay to meet a given performance requirement for the consensus approach.

Notation: We use boldface lowercase letters $\boldsymbol{x}$ for column vectors, and boldface uppercase letters $\boldsymbol{X}$ for matrices. In particular, $\boldsymbol{I}_M$ denotes an $M \times M$ identity matrix, $\mathbf{1}$ is the all-one vector, and $\mathbf{0}$ is the all-zero vector. Sets are described by caligraphic letters $\mathcal{X}$. The cardinality of a set $\mathcal{X}$ is denoted by $|\mathcal{X}|$. The transpose of a vector is denoted by $[\cdot]^{\mathsf{T}}$. Expectation of a random variable is denoted by $\mathbb{E}\{\cdot\}$, and $\mathrm{Var}(\cdot)$ indicates the variance of a random variable or covariance matrix of a random vector. The indicator function of a set $\mathcal{X}$ is written as $\mathbb{I}_{\mathcal{X}}(\cdot)$.

II. System Model

A. Sensor and Network Model

The system model is illustrated in Fig. 1. We consider a one-dimensional network of N nodes placed equally spaced on a straight line. Without loss of generality we set the coordinate of sensor k, $k = 0, \ldots, N-1$, to $s_k = k$. The sensors measure the intensity of a real-valued signal $x(s)$ in their respective

This research was supported, in part, by the European Research Council, under Grant No. 258418 (COOPNET), and the Swedish Research Council, under Grant No. 2011-5961.

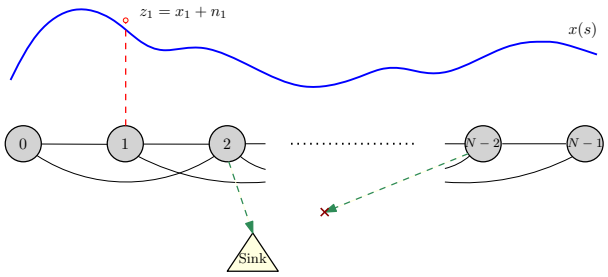

Fig. 1. Illustration of the system model. The grey circles are the sensor nodes measuring the intensity of the signal $x(s)$. They perform local processing using sensor-to-sensor communication (black lines). Sensors 2 and $N-2$ send packets to the sink, but the packet from sensor $N-2$ is erased.

spatial coordinates. The observation of sensor k is

$$z_k = x_k + n_k, \tag{1}$$

where $x_k = x(k)$ and the n_k's are spatially white Gaussian noise samples with variance σ_n^2. The observations are stacked in a vector $\boldsymbol{z} = [z_0, \ldots, z_{N-1}]^\mathsf{T}$.

Each node can communicate over a licensed spectrum with a base station serving as a sink, or fusion center (FC), incurring a fixed cost (bandwidth usage) C. The node-to-sink links are modeled as independent erasure channels with packet erasure probability p.[1] Communication from the nodes to the FC relies on orthogonal channels, and thus incurs no delay.

The nodes can also exchange information locally with nearby nodes using broadcasting over a shared (unlicensed) channel. To avoid packet collisions the transmissions are scheduled using a spatial time division multiple access (S-TDMA) protocol. Each node is allowed to transmit only in an assigned time slot, which is reused by other nodes if they are spatially far apart. Therefore, the local communication will incur a delay D (expressed in a number of TDMA slots), but is on the other hand considered to be cost-free. We use a disc model with radius R to determine if two nodes are connected. For later use, we denote by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ the undirected graph describing the network, where $\mathcal{V}$ is the set of nodes and $\mathcal{E}$ the set of edges connecting the nodes.

B. Signal Model

We consider a smooth, band-limited spatial signal $x(s)$, sampled as $\boldsymbol{x} = [x_0, \ldots, x_{N-1}]^\mathsf{T}$ with energy $E_X = \sum_{k=0}^{N-1} |x_k|^2$. Furthermore, we assume that there exists a transformation $\boldsymbol{\theta} = \boldsymbol{T}\boldsymbol{x}$ such that $\boldsymbol{\theta} \in \mathbb{C}^{N\times 1}$ is K-sparse, i.e., $\boldsymbol{\theta}$ has $K \ll N$ nonzero elements. In our case, the signal $x(s)$ is regarded as sparse in the spatial frequency domain, owing to the smoothness of $x(s)$. Since nodes are equispaced, we can use a discrete Fourier transform (DFT) matrix as $\boldsymbol{T}$, with entries

$$T_{ml} = \frac{1}{\sqrt{N}} \exp\left(-j\frac{2\pi ml}{N}\right), \tag{2}$$

for $m, l \in \{0, \ldots, N-1\}$. The entries of $\boldsymbol{\theta}$ are then the sampled spatial frequencies of $x(s)$. We will denote the average signal-to-noise ratio per sample $E_X/(N\sigma_n^2)$ as SNR.

C. Goal

Given the observations $\boldsymbol{z}$ and the system model outlined above, the goal is to reconstruct $\boldsymbol{x}$ at the sink such that a certain reconstruction error is guaranteed.

III. Compressed Sensing Background

A. Definition and Performance Measure

Let $\boldsymbol{x}$ and $\boldsymbol{T}\boldsymbol{x} = \boldsymbol{\theta}$ be as described in Section II-B. Also, let $\boldsymbol{A} \in \mathbb{R}^{M\times N}$ be a measurement matrix where $M \ll N$, and define the compression

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{z} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e}_{\text{obs}}, \tag{3}$$

where $\boldsymbol{e}_{\text{obs}} = \boldsymbol{A}\boldsymbol{n}$, and $\boldsymbol{n} = [n_0, \ldots, n_{N-1}]^\mathsf{T}$. Since $M \ll N$, recovering $\boldsymbol{x}$ from $\boldsymbol{y}$ is an ill-posed problem, as there are infinitely many vectors $\tilde{\boldsymbol{x}}$ that satisfy $\boldsymbol{y} = \boldsymbol{A}\tilde{\boldsymbol{x}}$. However, we can exploit the knowledge about the sparsity of $\boldsymbol{x}$ in the transform domain. If $\boldsymbol{A}$ satisfies the restricted isometry property (RIP) [3], and $M \geq \rho K \log(N/K)$,[2] we can recover $\boldsymbol{x}$ from $\boldsymbol{y}$ by considering the following ℓ_1-norm minimization problem,

$$\text{minimize} \quad \|\boldsymbol{\theta}\|_1 \tag{4a}$$

$$\text{subject to} \quad \|\boldsymbol{A}\boldsymbol{T}^{-1}\boldsymbol{\theta} - \boldsymbol{y}\|_2 \leq \varepsilon, \tag{4b}$$

called BPDN [4]. If, for a given matrix $\boldsymbol{A}$, there exists a constant $\delta_K \in (0,1)$ such that the following inequality holds for all K-sparse vectors $\boldsymbol{v}$,

$$(1-\delta_K)\|\boldsymbol{v}\|_2^2 \leq \|\boldsymbol{A}\boldsymbol{v}\|_2^2 \leq (1+\delta_K)\|\boldsymbol{v}\|_2^2, \tag{5}$$

then $\boldsymbol{A}$ satisfies the RIP of order K. The computation of δ_K is NP-hard. In [3] and [13] it was shown that if $\boldsymbol{A}$ is a Gaussian random matrix with i.i.d. entries $A_{ml} \sim \mathcal{N}(0, 1/M)$, then $\boldsymbol{A}$ satisfies the RIP with very high probability.

Assuming $\delta_{2K} < \sqrt{2} - 1$ and $\varepsilon \geq \|\boldsymbol{e}_{\text{obs}}\|_2$, the ℓ_2-norm of the reconstruction error of BPDN is upper bounded by [4], [14]

$$\|\boldsymbol{x} - \boldsymbol{x}^\star\|_2 \leq \frac{C_0}{\sqrt{K}}\|\boldsymbol{T}^{-1}(\boldsymbol{\theta} - \boldsymbol{\theta}_K)\|_1 + C_1\varepsilon, \tag{6}$$

where $\boldsymbol{x}^\star = \boldsymbol{T}^{-1}\boldsymbol{\theta}^\star$, in which $\boldsymbol{\theta}^\star$ is the solution to (4), $\boldsymbol{\theta}_K$ is the best K-sparse approximation of the transformed underlying signal, and $C_0, C_1 \geq 0$ are constants that depend on δ_K [4]. Here, we only consider strictly K-sparse signals, meaning there are at most K non-zero components in $\boldsymbol{\theta}$. Hence, the first term on the right hand side of (6) is zero.

Since the entries of $\boldsymbol{A}$ are i.i.d. Gaussian, it follows that $\boldsymbol{e}_{\text{obs}} \sim \mathcal{N}\left(0, (N/M)\sigma_n^2\boldsymbol{I}_M\right)$, so that $\|\boldsymbol{e}_{\text{obs}}\|_2$ is distributed according to a scaled χ_M-distribution. Hence, by Taylor series expansion, $\mathbb{E}_{\boldsymbol{A},\boldsymbol{n}}\{\|\boldsymbol{e}_{\text{obs}}\|_2\} \approx \sigma_n\sqrt{N}(1-1/(4M))$ and $\text{Var}_{\boldsymbol{A},\boldsymbol{n}}(\|\boldsymbol{e}_{\text{obs}}\|_2) \approx (N/M)\sigma_n^2(1/2 - 1/(8M))$. Therefore, to

[1] The event that all packets sent to the sink are erased is not considered.

[2] The parameter $\rho \in (0,1)$ is independent of M, N, and K. An exact expression can be found in [12].

satisfy $\varepsilon \geq \|\boldsymbol{e}_{\text{obs}}\|_2$ with high probability, ε should be choosen as

$$\varepsilon = \varepsilon_{\text{ref}} = \sigma_n\sqrt{N}\left(\left(1-\frac{1}{4M}\right)+\lambda\sqrt{\frac{1}{2M}-\frac{1}{8M^2}}\right) \quad (7)$$

$$\approx \sigma_n\sqrt{N}\left(1+\lambda\sqrt{\frac{1}{2M}}\right), \quad (8)$$

where $\lambda \geq 0$ is used to achieve a desired confidence level.

B. Distributed Compressed Sensing for Networked Data

We observe that the compression in (3) can be written as a sum of linear projections of the measurements z_k onto the corresponding column $\boldsymbol{a}_k$ of $\boldsymbol{A}$,

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{z} = \sum_{k=0}^{N-1} \boldsymbol{a}_k z_k. \quad (9)$$

If we generate $\boldsymbol{a}_k$ in sensor k, it can compute its contribution $\boldsymbol{w}_k = \boldsymbol{a}_k z_k$ to the compression. By distributing the local projections $\boldsymbol{w}_k$ in the network using sensor-to-sensor communication and local processing in the sensors, we can compute (9) in a decentralized manner. Consequently, this compression reduces the number of sensors that need to convey information to the sink, effectively reducing the communication cost at the expense of a delay induced by the local communication. In Sections IV and V we present two approaches to such distributed processing for which we determine the node-to-sink communication cost, inter-node communication delay, and an upper bound on the reconstruction error.

IV. Distributed Linear Projections using Clustering

A. Cluster Formation and Operation

A set of nodes $\mathcal{L} \subseteq \mathcal{V}$, $|\mathcal{L}| = L$, is selected to act as aggregating nodes (clusterheads), such that clusterhead $i \in \mathcal{L}$ collects information from a subset (cluster) $\mathcal{C}_i \subseteq \mathcal{V}$, of the sensors in the network. The clusterhead selection is done with respect to the local communication range such that each clusterhead is located at the center of its cluster, which has radius R. The clusters are disjoint, i.e., $\mathcal{C}_i \cap \mathcal{C}_l = \emptyset$ for $i \neq l$, and $\bigcup_{i\in\mathcal{L}} \mathcal{C}_i = \mathcal{V}$. Note that depending on N and R, one of the clusters at the boundary may be smaller than the others. The number of clusters is given by

$$L = \left\lceil \frac{N}{2R+1} \right\rceil. \quad (10)$$

Node k computes its local linear projection $\boldsymbol{w}_k = \boldsymbol{a}_k z_k$ and sends it to its clusterhead. Clusterhead i computes

$$\boldsymbol{y}_{\mathcal{C}_i} = \sum_{k\in\mathcal{C}_i} \boldsymbol{w}_k. \quad (11)$$

Finally, the clusterheads transmit their partial information to the sink. Since the clusters are disjoint, the sink computes $\boldsymbol{y} = \boldsymbol{y}_{\mathcal{C}_1} + \cdots + \boldsymbol{y}_{\mathcal{C}_L} = \boldsymbol{A}\boldsymbol{z}$, and reconstructs $\boldsymbol{x}$ using BPDN.

B. Cost and Delay

The total communication cost is $C_{\text{tot}} = CL$. The delay is given by the number of time slots in the S-TDMA needed to schedule a broadcast transmission for every non-clusterhead node. Due to the cluster formation and communication model that we consider, there is no interference from nodes in a cluster to the neighboring clusterheads. Hence, the delay D is given by the maximum node degree of the clusterheads,

$$D = \begin{cases} 2R & 0 \leq 2R < N \\ N-1 & 2R \geq N. \end{cases} \quad (12)$$

C. Reconstruction Performance and Robustness

Define the set $\mathcal{D} \subseteq \mathcal{L}$ as the set of clusterheads whose packets are erased during transmission to the sink. The sink is assumed to have no knowledge of $\mathcal{D}$, but attempts to recover $\boldsymbol{x}$ assuming it has received the correct compression $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{z}$. The resulting compression at the sink given a set of packet erasures $\mathcal{D}$ is

$$\tilde{\boldsymbol{y}} = \sum_{i\in\mathcal{L}\setminus\mathcal{D}} \boldsymbol{y}_{\mathcal{C}_i} = \boldsymbol{A}\boldsymbol{x} + (\boldsymbol{A}-\boldsymbol{B})\boldsymbol{n} - \boldsymbol{B}\boldsymbol{x} = \boldsymbol{y} - \boldsymbol{B}\boldsymbol{z}, \quad (13)$$

where $\boldsymbol{B}$ is a matrix whose nonzero columns, corresponding to the nodes whose clusterhead packet was erased, are equal to the corresponding columns of $\boldsymbol{A}$. Therefore, for packet erasure probability $p = 0$, we have $\boldsymbol{B} = \boldsymbol{0}$ and $\tilde{\boldsymbol{y}} = \boldsymbol{y}$, while for $p \neq 0$ we have to account for $\boldsymbol{B}\boldsymbol{z}$ when setting ε in the BPDN for (6) to hold. The following Theorem describes how ε should be selected.

Theorem 1. *Given the model described in Section II, $\boldsymbol{A}$ as described in Section III-A, and the compression in (13), the choice of ε that guarantees a stable recovery using BPDN is*

$$\varepsilon = \varepsilon_{\text{ref}}\sqrt{1-\left(1-\frac{1-p}{1-p^{\lceil N/(2R+1)\rceil}}\right)(1-\mathsf{SNR})}. \quad (14)$$

The proof is given in Appendix A.

V. Distributed Linear Projections using Consensus

An alternative approach is to compute $\boldsymbol{y}$ from (9) directly in the network by using a fully distributed algorithm. Here, we propose the use of average consensus.

A. The Consensus Algorithm

We can express (9) as

$$\boldsymbol{y} = \sum_{k=0}^{N-1} \boldsymbol{w}_k = N\left(\frac{1}{N}\sum_{k=0}^{N-1} \boldsymbol{w}_k\right) = N\bar{\boldsymbol{w}}. \quad (15)$$

We use average consensus to estimate $\bar{\boldsymbol{w}}$ in the network. The estimate is then used at the sink to compute (15). Let $\boldsymbol{w}_k(0) = \boldsymbol{w}_k$ be the initial value at sensor k. The updating rule of average consensus [15] is given by

$$\boldsymbol{w}_k(i) = \boldsymbol{w}_k(i-1) + \xi \sum_{v\in\mathcal{M}_k} \left(\boldsymbol{w}_k(i-1) - \boldsymbol{w}_v(i-1)\right), \quad (16)$$

where $\mathcal{M}_k \subset \mathcal{V}$ is the set of neighboring sensors of sensor k, ξ is the algorithm step size, and i is the iteration index. We can also express (16) in matrix form as

$$\boldsymbol{W}(i) = \boldsymbol{P}\boldsymbol{W}(i-1) = \boldsymbol{P}^i\boldsymbol{W}(0), \tag{17}$$

where $\boldsymbol{W}(0) = [\boldsymbol{w}_1(0), \ldots, \boldsymbol{w}_N(0)]^\mathsf{T}$ and $\boldsymbol{P} = \boldsymbol{I}_N - \xi\boldsymbol{L}$, in which $\boldsymbol{L}$ denotes the graph Laplacian of $\mathcal{G}$. By properties of the consensus algorithm [15], $\bar{\boldsymbol{w}}$ is conserved in each iteration,

$$\bar{\boldsymbol{w}} = \frac{1}{N}\sum_{k=0}^{N-1}\boldsymbol{w}_k(0) = \frac{1}{N}\sum_{k=0}^{N-1}\boldsymbol{w}_k(i), \tag{18}$$

irrespective of i. If ξ is chosen small enough [15], the algorithm is monotonically converging in the limit $i \to +\infty$ to the average in all sensor nodes,

$$\lim_{i\to\infty}\boldsymbol{w}_k(i) = \bar{\boldsymbol{w}}. \tag{19}$$

After a certain number of iterations I, a set $\mathcal{L} \subseteq \mathcal{V}$, with $|\mathcal{L}| = L$, of randomly chosen sensors communicate their estimates $\boldsymbol{w}_k(I)$ of $\bar{\boldsymbol{w}}$ to the sink. However, due to erasures, the sink estimates $\boldsymbol{y}$ from a set of nonerased packets $\tilde{\mathcal{L}} \subseteq \mathcal{L}$, $|\tilde{\mathcal{L}}| = \tilde{L}$, as

$$\hat{\boldsymbol{y}} = N\left(\frac{1}{\tilde{L}}\sum_{k\in\tilde{\mathcal{L}}}\boldsymbol{w}_k(I)\right). \tag{20}$$

Finally, $\hat{\boldsymbol{y}}$ (cf. (15)) is used in (4) to reconstruct $\boldsymbol{x}$.

B. Cost and Delay

As for clustering, the total communication cost is $C_{\text{tot}} = CL$. The delay is given by $D_{\text{cons}} = DI$, where D was defined in (12).

C. Reconstruction Performance and Robustness

Due to the fact that average consensus only converges in the limit $i \to +\infty$, for any finite I there will be an error in each sensor estimate $\boldsymbol{w}_k(I)$ with respect to the true average $\bar{\boldsymbol{w}}$. The transmitted packets to the sink can also be erased. This results in a mismatch $\boldsymbol{e}_{\text{cons}} = \boldsymbol{y} - \hat{\boldsymbol{y}}$ between the desired compression and the compression calculated using average consensus. As for the clustering case, in order to guarantee that the reconstruction error is upper bounded by (6), this perturbation has to be accounted for when setting ε in the BPDN. The following Theorem states how this should be done.

Theorem 2. *Let μ_2 be the second largest eigenvalue of $\boldsymbol{P}$. Given the model in Section II, $\boldsymbol{A}$ as described in Section III-A, and the compression in (20), the choice of ε that guarantees a stable recovery using BPDN is*

$$\varepsilon = \varepsilon_{\text{ref}}\left(1 + \mu_2^I\sqrt{(1+\mathsf{SNR})\,\Phi}\right), \tag{21}$$

where[3]

$$\Phi = \frac{N(1-p^L)}{L(1-p)}\left(\frac{N - \frac{L(1-p)}{(1-p^L)}}{N-1}\right). \tag{22}$$

The proof is given in Appendix B.

VI. Results and Discussion

In this section, we evaluate the cost-delay tradeoff of the clustering and consensus approaches, i.e., how the reconstruction error scales with the number of iterations I and the number of nodes transmitting to the sink L, and compare the robustness to packet erasures. We fix $N = 100$, $M = 20$, $R = 10$, $\sigma_n^2 = 0.01$, and $E_X = 3$, giving $\mathsf{SNR} = 3$ in linear scale. The figures are created by computing ε using the expressions in Theorems 1 and 2, where the upper bound $\mu_2 \leq \cos(\pi R/2N)$ is used in (21), assuming ξ is chosen optimally [15]. Since C_1 in (6) is NP-hard to compute, we normalize the error with respect to C_1. Also, since M, N, and σ_n^2 are fixed, we also normalize with respect to ε_{ref}. Hence, the normalized error is equal to $\zeta_{\text{norm}} = \|\boldsymbol{x} - \boldsymbol{x}^\star\|_2/(C_1\varepsilon_{\text{ref}})$. Note that $\zeta_{\text{norm}} \geq 1$.

A. Cost-Delay Tradeoff

Fig. 2 shows the boundaries of the regions giving a normalized error lower than the threshold $\nu = 1.1$ for packet erasure probabilities $p \in \{0, 0.05, 0.2, 0.5\}$. As can be seen, a higher packet erasure probability results in a boundary receding towards the top right corner, meaning that higher I and L are needed to meet ν. An important observation is also that the normalized error ζ_{norm} is nonincreasing in I and L. Looking at the slope of the curves, we see that there are differences in how much delay we must tolerate in order to lower communication cost. For example when $p = 0$, for low and high costs, we need to increase delay significantly, while

[3]The value of ε in (21) may be very conservative, since the upper bound on the convergence rate of consensus (see (47) in Appendix B) may be very loose. Consider the eigenvalue decomposition of $\boldsymbol{P} = \boldsymbol{Q}\boldsymbol{M}\boldsymbol{Q}^{-1}$, and let $\boldsymbol{\alpha} = \boldsymbol{Q}^{-1}\boldsymbol{w}$ be the projection of the data $\boldsymbol{w}$ onto the eigenspace of $\boldsymbol{P}$. Order the eigenvalues of $\boldsymbol{P}$ as $1 = \mu_1 > \mu_2 \geq \cdots \geq \mu_N$. Then, the disagreement after I iterations of consensus on $\boldsymbol{w}$ is given by

$$\|\boldsymbol{w}(I) - \bar{\boldsymbol{w}}\|_2^2 = \sum_{k=2}^{N}\mu_k^{2I}\alpha_k^2, \tag{23}$$

which can be upper bounded by

$$\sum_{k=2}^{N}\mu_k^{2I}\alpha_k^2 \leq \mu_2^{2I}\sum_{k=2}^{N}\alpha_k^2 \overset{\text{(a)}}{=} \mu_2^{2I}\|\boldsymbol{w}(0) - \bar{\boldsymbol{w}}\|_2^2, \tag{24}$$

where (a) follows since α_1 is the entry that corresponds to the eigenvector of μ_1, and thus the initial disagreement is $\|\boldsymbol{w}(0) - \bar{\boldsymbol{w}}\|_2^2 = \sum_{k=2}^{N}\alpha_k^2$. In general, (23) is hard to compute and (24) may be loose. If the support of $\boldsymbol{\alpha}$ is concentrated to those entries corresponding to the smaller eigenvalues, $\|\boldsymbol{w}(I) - \bar{\boldsymbol{w}}\|_2^2$ decreases much faster than μ_2^{2I} in the first iterations. However, after enough iterations the smaller eigenvalues have diminished, and the convergence rate is dominated by μ_2. In our case, the data $\boldsymbol{w}$ is Gaussian, and since the columns of $\boldsymbol{Q}^{-1}$ form an orthonormal basis in $\mathbb{R}^N$, $\boldsymbol{\alpha}$ is also Gaussian with the same mean and variance. Therefore, the power of $\boldsymbol{\alpha}$ is spread evenly in its entries. Consequently, the bound is loose for our signals and consensus behaves much better with respect to I than shown in Figs. 2 and 3.

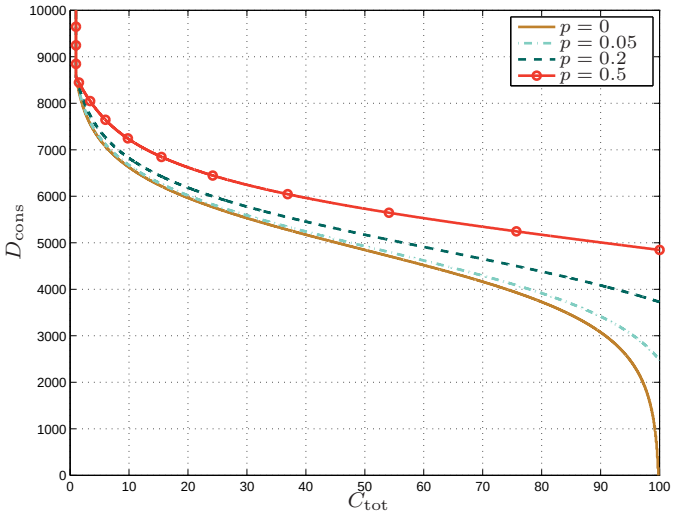

Fig. 2. The figure shows the boundaries of systems satisfying the given error threshold ν. The area above the graphs are the regions of points $(C_{\text{tot}}, D_{\text{cons}})$ satisfying $\zeta_{\text{norm}} \leq \nu$, where $\nu = 1.1$, for $p \in \{0, 0.05, 0.2, 0.5\}$, when using average consensus with $R = 10$.

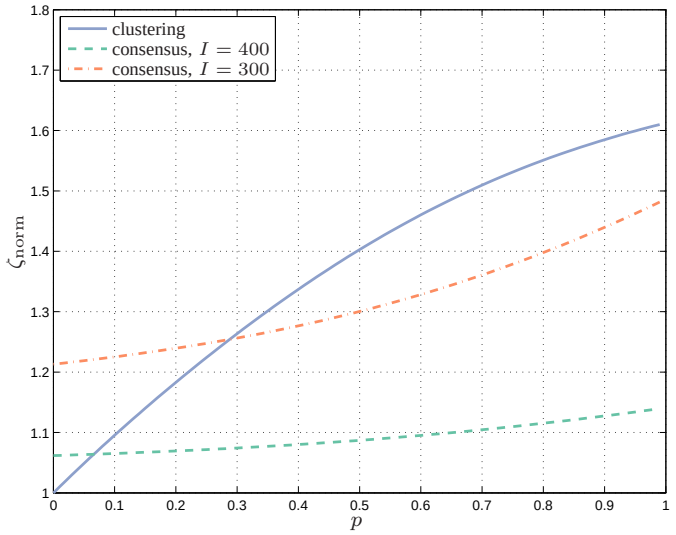

Fig. 3. This plot shows ζ_{norm} for different packet erasure probabilities of clustering and two cases of consensus with $I \in \{300, 400\}$, where $R = 10$ and $L = \lceil N/(2R+1) \rceil$

for medium costs the curves are flatter and a smaller increase in delay is sufficient to reduce cost.

For clustering, D and C_{tot} are implicitly given by R through (12) and (10). Hence, there is no tradeoff as such for the clustering. The implication is that for larger p we cannot increase cost or delay to ensure that $\zeta_{\text{norm}} \leq \nu$.

B. Robustness to Packet Erasures

Fig. 3 depicts the behavior of ζ_{norm} with respect to p. From the slope of the curves we see that consensus is less sensitive to packet erasures as compared to clustering. This is in line with the results in Theorems 1 and 2, where $\zeta_{\text{norm}} \propto \sqrt{p}$ and $\zeta_{\text{norm}} \propto 1/\sqrt{1-p}$, for clustering and consensus, respectively (see (14) and (21)). Note that the source of error is different for clustering and consensus. Both approaches are affected by packet erasures, but in different manners. For clustering, if an erasure occurs, that information is lost, while for consensus the estimation step (20) at the sink is affected only to a small degree. This is because the consensus algorithm disseminates the information throughout the network, making it more robust to packet erasures. On the other hand, for consensus ζ_{norm} is dominated by the disagreement between the estimates at the nodes and the true average. This explains the superiority of clustering for small p. However, the disagreement decreases exponentially in I, so ζ_{norm} can be made arbitrarily small by increasing I.

VII. Conclusion

We derived closed-form expressions for the upper bound on the ℓ_2-norm of the reconstruction error for a clustering and a consensus approach to distributed compressed sensing in WSNs. For the consensus approach, the expression can be used to trade off cost and delay such that the reconstruction error is guaranteed to satisfy a given performance requirement with high probability. We also analyzed the robustness to erasures of packets sent to the sink. If a large enough number of iterations is allowed, consensus is more robust than clustering, except for very small packet erasure probabilities. Moreover, by increasing the number of iterations, the additional error caused by the consensus algorithm and packet erasures can be made arbitrarily small. Another benefit of the consensus is that there is no need to form clusters, which can be a hard task, especially if the sensors are mobile. Future research includes unreliabe sensor-to-sensor communication, uncertainty in the position of the nodes, and more general network topologies.

Appendix A
Proof of Theorem 1

When using clustering, the compressed vector received by the sink is

$$\tilde{\boldsymbol{y}} = \boldsymbol{A}\boldsymbol{x} + (\boldsymbol{A} - \boldsymbol{B})\boldsymbol{n} - \boldsymbol{B}\boldsymbol{x}. \tag{25}$$

Define $\boldsymbol{C} = \boldsymbol{A} - \boldsymbol{B}$, and the total perturbation $\boldsymbol{u} = \boldsymbol{C}\boldsymbol{n} - \boldsymbol{B}\boldsymbol{x}$. Let $\boldsymbol{c}_k$ and $\boldsymbol{b}_k$ be the kth column vector of $\boldsymbol{C}$ and $\boldsymbol{B}$, respectively. Then, for each node k we have $\boldsymbol{u}_k = \boldsymbol{c}_k n_k - \boldsymbol{b}_k x_k$, and $\boldsymbol{u} = \sum_{k=0}^{N-1} \boldsymbol{u}_k$. Therefore, for (6) to hold, we need $\varepsilon \geq \|\boldsymbol{u}\|_2$. Denote by $\mathcal{H}$ the set of nodes whose information is not erased, i.e., $\mathcal{H} = \{k \ : \ k \in \mathcal{C}_j, j \notin \mathcal{D}\}$, where $|\mathcal{H}| = H$. Note that H is a zero-truncated binomial random variable with parameters L and p. It is easy to see that $\mathbb{E}_{\mathcal{H},\boldsymbol{A},\boldsymbol{n}}\{\boldsymbol{u}_k\} = \boldsymbol{0}$, hence the covariance matrix is

$$\begin{aligned}\mathbb{E}_{\mathcal{H},\boldsymbol{A},\boldsymbol{n}}\left\{\boldsymbol{u}_k\boldsymbol{u}_k^{\mathsf{T}}\right\} &= \mathbb{E}_{\mathcal{H},\boldsymbol{A}}\left\{\boldsymbol{c}_k\boldsymbol{c}_k^{\mathsf{T}}\right\}\mathbb{E}_{\boldsymbol{n}}\left\{n_k^2\right\} \\ &\quad + \mathbb{E}_{\mathcal{H},\boldsymbol{A}}\left\{\boldsymbol{b}_k\boldsymbol{b}_k^{\mathsf{T}}\right\}x_k^2. \end{aligned} \tag{26}$$

For notational convenience, we drop the subscript indicating over which variable the expectation is taken. We observe that

for $k \notin \mathcal{H}$, $\boldsymbol{c}_k = \mathbf{0}$ and $\boldsymbol{b}_k = \boldsymbol{a}_k$, and for $k \in \mathcal{H}$, $\boldsymbol{c}_k = \boldsymbol{a}_k$ and $\boldsymbol{b}_k = \mathbf{0}$, thus

$$\mathbb{E}\left\{\boldsymbol{u}_k\boldsymbol{u}_k^{\mathsf{T}}\right\} = \begin{cases} \mathbb{E}\left\{\boldsymbol{c}_k\boldsymbol{c}_k^{\mathsf{T}}\right\}\mathbb{E}\left\{n_k^2\right\} = \frac{\sigma_n^2}{M}\boldsymbol{I}_M & k \in \mathcal{H} \\ \mathbb{E}\left\{\boldsymbol{b}_k\boldsymbol{b}_k^{\mathsf{T}}\right\} x_k^2 = \frac{x_k^2}{M}\boldsymbol{I}_M & k \notin \mathcal{H}. \end{cases} \tag{27}$$

It follows that

$$\mathbb{E}\left\{\boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}\right\} = \mathbb{E}\left\{\left(\sum_{k=0}^{N-1}\boldsymbol{u}_k\right)\left(\sum_{k=0}^{N-1}\boldsymbol{u}_k\right)^{\mathsf{T}}\right\} \tag{28}$$

$$\overset{(a)}{=} \mathbb{E}\left\{\sum_{k=0}^{N-1}\boldsymbol{u}_k\boldsymbol{u}_k^{\mathsf{T}}\right\} \tag{29}$$

$$= \sum_{k=0}^{N-1}\left(\mathbb{E}\left\{\boldsymbol{u}_k\boldsymbol{u}_k^{\mathsf{T}}\mathbb{I}_{\mathcal{H}}(k)\right\} + \mathbb{E}\left\{\boldsymbol{u}_k\boldsymbol{u}_k^{\mathsf{T}}\mathbb{I}_{\mathcal{V}\setminus\mathcal{H}}(k)\right\}\right), \tag{30}$$

where (a) follows since all $\boldsymbol{u}_k$'s are mutually independent. The probability that $k \notin \mathcal{H}$ is given by

$$p_{\mathcal{H}} = 1 - \frac{1-p}{1-p^{\lceil N/(2R+1)\rceil}}. \tag{31}$$

Then, we have

$$\mathbb{E}\left\{\boldsymbol{u}\boldsymbol{u}^{\mathsf{T}}\right\} = \frac{1}{M}\left(N(1-p_{\mathcal{H}})\sigma_n^2 + Np_{\mathcal{H}}\frac{E_X}{N}\right)\boldsymbol{I}_M \tag{32}$$

$$= \frac{\sigma_n^2 N}{M}\left(1 - p_{\mathcal{H}}(1-\mathsf{SNR})\right)\boldsymbol{I}_M \triangleq \sigma_{\boldsymbol{u}}^2\boldsymbol{I}_M. \tag{33}$$

For large enough N, $\boldsymbol{u} \sim \mathcal{N}(\mathbf{0}, \sigma_{\boldsymbol{u}}^2\boldsymbol{I}_M)$, and consequently $\|\boldsymbol{u}\|_2$ is distributed according to a scaled χ_M-distribution. Hence, $\mathbb{E}\{\|\boldsymbol{u}\|_2\} = \sigma_{\boldsymbol{u}}\sqrt{M}(1-1/4M)$ and $\mathrm{Var}(\|\boldsymbol{u}\|_2) = \sigma_{\boldsymbol{u}}^2(1/2 - 1/8M)$. Therefore, using $\sigma_{\boldsymbol{u}}^2$ as defined in (33), the robust choice for ε is

$$\varepsilon = \mathbb{E}\{\|\boldsymbol{u}\|_2\} + \lambda\sqrt{\mathrm{Var}(\|\boldsymbol{u}\|_2)} \tag{34}$$

$$= \varepsilon_{\mathrm{ref}}\sqrt{1 - \left(1 - \frac{1-p}{1-p^{\lceil N/(2R+1)\rceil}}\right)(1-\mathsf{SNR})}. \tag{35}$$

Appendix B
Proof of Theorem 2

The vector received by the sink using consensus is

$$\hat{\boldsymbol{y}} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e}_{\mathrm{obs}} + \boldsymbol{e}_{\mathrm{cons}}. \tag{36}$$

In order to guarantee stable reconstruction $\varepsilon \geq \|\boldsymbol{e}_{\mathrm{obs}} + \boldsymbol{e}_{\mathrm{cons}}\|_2$. By the triangle inequality, we have

$$\|\boldsymbol{e}_{\mathrm{obs}} + \boldsymbol{e}_{\mathrm{cons}}\|_2 \leq \|\boldsymbol{e}_{\mathrm{obs}}\|_2 + \|\boldsymbol{e}_{\mathrm{cons}}\|_2. \tag{37}$$

Thus, we choose $\varepsilon \geq \|\boldsymbol{e}_{\mathrm{obs}}\|_2 + \|\boldsymbol{e}_{\mathrm{cons}}\|_2$. The statistics of the first term on the right hand side of (37) are given in Section III. It remains to determine the contribution from the consensus. Since all dimensions of $\hat{\boldsymbol{y}}$ are i.i.d. we can calculate the statistics from one dimension and deduce what the total contribution is. We fix the number of iterations I, the number of queried nodes L, and the data $\boldsymbol{w}_k$. Define the disagreement between the estimate from $\tilde{L}$ received packets and the true average after I iterations for each dimension $m = 1, \ldots, M$ as

$$\Delta_m(I, L) = \frac{1}{\tilde{L}}\left(\sum_{k\in\tilde{\mathcal{L}}} w_{k,m}(I)\right) - \bar{w}_m = \hat{w}_m(I, \tilde{L}) - \bar{w}_m, \tag{38}$$

where $w_{k,m}(I)$ is the mth element of the vector $\boldsymbol{w}_k(I)$, $\bar{w}_m$ is the average over the mth dimension, and $\hat{w}_m(I, \tilde{L})$ is the estimate of $\bar{w}_m$. For notational convenience we drop the subscript indicating the dimension, and the dependencies on I and L. Now, there are two sources of randomness: (i) the set of queried nodes $\mathcal{L} \subseteq \mathcal{V}$, which is randomly selected; (ii) the number of nonerased packets $\tilde{L} \leq L$, due to random packet erasures. Since w_k's are fixed, $\bar{w}$ is constant. Hence,

$$\mathrm{Var}_{\tilde{L},\mathcal{L}}(\Delta) = \mathrm{Var}_{\tilde{L},\mathcal{L}}(\hat{w}) = \mathbb{E}_{\tilde{L},\mathcal{L}}\left\{(\hat{w} - \bar{w})^2\right\} \tag{39}$$

$$= \mathbb{E}_{\tilde{L}}\left\{\mathbb{E}_{\mathcal{L}|\tilde{L}}\left\{(\hat{w} - \bar{w})^2\right\}\right\}. \tag{40}$$

The estimate $\hat{w}$ is an estimate by simple random sampling from a finite population of size N. Then

$$\mathbb{E}_{\tilde{L}}\left\{\mathbb{E}_{\mathcal{L}|\tilde{L}}\left\{(\hat{w} - \bar{w})^2\right\}\right\} \tag{41}$$

$$= \mathbb{E}_{\tilde{L}}\left\{\frac{1}{\tilde{L}N}\left(\sum_{k=0}^{N-1}(w_k - \bar{w})^2\right)\frac{N-\tilde{L}}{N-1}\right\} \tag{42}$$

$$= \frac{1}{N(N-1)}\left(\sum_{k=0}^{N-1}(w_k - \bar{w})^2\right)\left(N\mathbb{E}_{\tilde{L}}\left\{\frac{1}{\tilde{L}}\right\} - 1\right), \tag{43}$$

where the first equality is due to [16, Thm. 7.3.1B]. The expectation of the inverse of $\tilde{L}$ is [17]

$$\mathbb{E}_{\tilde{L}}\left\{\frac{1}{\tilde{L}}\right\} \approx (1-p^L)\frac{1}{L(1-p)} \triangleq \frac{1}{\bar{L}}. \tag{44}$$

If we consider again the dependence of w_k on I and let $\boldsymbol{w}(I) = [w_1(I), \ldots, w_N(I)]^{\mathsf{T}}$, we have

$$\mathrm{Var}_{\tilde{L},\mathcal{L}}(\Delta) = \frac{1}{\bar{L}}\left(\frac{N-\bar{L}}{N(N-1)}\right)\|\boldsymbol{w}(I) - \bar{w}\mathbf{1}\|_2^2. \tag{45}$$

The convergence rate of the ℓ_2-norm $\|\boldsymbol{w}(I) - \bar{w}\mathbf{1}\|_2$ is defined as

$$\varrho = \lim_{I\to\infty}\left(\frac{\|\boldsymbol{w}(I) - \bar{w}\mathbf{1}\|_2}{\|\boldsymbol{w}(0) - \bar{w}\mathbf{1}\|_2}\right)^{1/I}, \quad \boldsymbol{w}(0) \neq \bar{w}\mathbf{1}, \tag{46}$$

which can be upper bounded by $\varrho \leq \mu_2^I$ [15]. Consequently, we have

$$\|\boldsymbol{w}(I) - \bar{w}\mathbf{1}\|_2^2 \leq \|\boldsymbol{w}(0) - \bar{w}\mathbf{1}\|_2^2\mu_2^{2I}. \tag{47}$$

Now, considering the randomness of $\boldsymbol{A}$ and $\boldsymbol{n}$,

$$\mathrm{Var}_{\tilde{L},\mathcal{L},\boldsymbol{A},\boldsymbol{n}}(\Delta) \leq \frac{1}{\bar{L}}\left(\frac{N-\bar{L}}{N(N-1)}\right)\mathbb{E}_{\boldsymbol{A},\boldsymbol{n}}\left\{\|\boldsymbol{w}(0) - \bar{w}\mathbf{1}\|_2^2\right\}\mu_2^{2I}. \tag{48}$$

Furthermore

$$\mathbb{E}_{\boldsymbol{A},\boldsymbol{n}}\left\{\|\boldsymbol{w}(0)-\bar{w}\mathbf{1}\|_2^2\right\}=\sum_{k=0}^{N-1}\mathbb{E}_{\boldsymbol{A},\boldsymbol{n}}\left\{(w_k(0)-\bar{w})^2\right\} \quad (49)$$

$$=\frac{N-1}{N}\left(\sum_{k=0}^{N-1}\frac{\sigma_n^2}{M}+\sum_{k=0}^{N-1}\frac{x_k^2}{M}\right) \quad (50)$$

$$\approx\frac{N\sigma_n^2+E_X}{M}, \quad (51)$$

where the last step follows since we consider very large N. Finally, we have

$$\sigma_\Delta^2\triangleq\mathrm{Var}_{\tilde{L},\mathcal{L},\boldsymbol{A},\boldsymbol{n}}(\Delta)\leq\frac{1}{\bar{L}}\left(\frac{N-\bar{L}}{N-1}\right)\frac{E_X+N\sigma_n^2}{NM}\mu_2^{2I}. \quad (52)$$

Due to the multiplication by N in (20), and since all dimensions m are i.i.d., $\boldsymbol{e}_{\mathrm{cons}}\sim\mathcal{N}(0,N^2\sigma_\Delta^2\boldsymbol{I}_M)$, hence $\|\boldsymbol{e}_{\mathrm{cons}}\|_2$ is distributed according to a scaled χ_M-distribution with $\mathbb{E}\{\|\boldsymbol{e}_{\mathrm{cons}}\|_2\}=N\sigma_\Delta\sqrt{M}(1-1/4M)$ and $\mathrm{Var}\left(\|\boldsymbol{e}_{\mathrm{cons}}\|_2\right)=N^2\sigma_\Delta^2(1/2-1/8M)$. Using (37) and the same argument as in (34), the robust choice of ε is

$$\varepsilon=\varepsilon_{\mathrm{ref}}\left(1+\mu_2^I\sqrt{(1+\mathsf{SNR})\,\Phi}\right), \quad (53)$$

where

$$\Phi=\frac{N(1-p^L)}{L(1-p)}\left(\frac{N-\frac{L(1-p)}{(1-p^L)}}{N-1}\right). \quad (54)$$

REFERENCES

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, Mar. 2002.

[2] M. F. Duarte, G. Shen, A. Ortega, and R. G. Baraniuk, "Signal compression in wireless sensor networks," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 118–135, 2012.

[3] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inf. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[4] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, Mar. 2006.

[5] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[6] H. Huang, S. Misra, W. Tang, H. Barani, and H. Al-Azzawi, "Applications of compressed sensing in communications networks," 2013. [Online]. Available: http://arxiv.org/abs/1305.3002

[7] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressive wireless sensing," in *Proc. 5th Int. Conf. on Information Processing in Sensor Networks*, Apr. 2006, pp. 134–142.

[8] J. Haupt, W. Bajwa, M. Rabbat, and R. Nowak, "Compressed sensing for networked data," *IEEE Sig. Proc. Mag.*, vol. 25, no. 2, pp. 92–101, Mar. 2008.

[9] C. Luo, F. Wu, J. Sun, and C. W. Chen, "Compressive data gathering for large-scale wireless sensor networks," in *Proc. 15th Annual Int. Conf. on Mobile computing and Networking*. ACM, Sep. 2009, pp. 145–156.

[10] H. Zheng, S. Xiao, X. Wang, and X. Tian, "On the capacity and delay of data gathering with compressive sensing in wireless sensor networks," in *Proc. IEEE Global Communications Conf.*, Dec. 2011.

[11] S. Patterson, Y. C. Eldar, and I. Keidar, "Distributed compressed sensing for static and time-varying networks," 2013. [Online]. Available: http://arxiv.org/abs/1308.6086/

[12] M. A. Davenport, M. F. Duarte, Y. C. Eldar, and G. Kutyniok, "Introduction to compressed sensing," in *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.

[13] R. Baraniuk, M. Davenport, R. Devore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, Dec. 2008.

[14] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Académie des Sciences*, vol. I, no. 346, pp. 589–592, May 2008.

[15] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, Jan. 2007.

[16] J. Rice, *Mathematical statistics and data analysis*. Cengage Learning, 2006.

[17] G. A. Rempała, "Asymptotic factorial powers expansions for binomial and negative binomial reciprocals," *Proceedings of the American Mathematical Society*, pp. 261–272, 2004.